

1

2 **Field of the Invention**

3 The invention relates to data communications and, more particularly, to improvements in
4 traffic shaping apparatus for packet-based data communications networks.

5 **Background Description**

6 Modern high speed packet-based data communications networks, such as those employing
7 Asynchronous Transfer Mode (ATM) technology, allow bandwidth to be reserved in advance. This
8 is useful, in particular, for the transport of real time data such as voice or video. In ATM networks
9 the requesting user must call the network control entity to set up a logical connection. Parameters
10 are given in the call set up request which specify the required bandwidth. If all the requested
11 parameters can be granted by the network controller the call is set up, otherwise it is refused.

12 An important problem in such networks is to ensure that traffic flowing in a particular
13 network connection respects the bandwidth limits which were negotiated for that call at set up time.
14 This is to guarantee that each connection only uses its share of the network resources, thereby
15 minimizing the probability of lost data or other problems resulting from congestion in the network.

16 Data loss within the network must be avoided not only because missing information can alter
17 the nature of the traffic -for example a missing cell can generate a glitch on a voice connection - but
18 also because data loss can provoke data retransmission by higher level communications protocols
19 thereby increasing the traffic on the connection and worsening network congestion - leading to a
20 higher probability of further data loss.

21 To ensure that traffic within a connection conforms to its negotiated bandwidth parameters,
22 traffic shaping or scheduling techniques have been developed for use at various points within the

1 network in order to modify the flow of data within a connection to conform to its agreed parameters.
2 Various policing schemes are also used in order to control and, if necessary, penalize traffic in the
3 network by discarding cells from connections which try to exceed their allotted resources.

4 A wide variety of traffic shaping and policing methods have been proposed and studied.
5 Many of these require a complex computation of an optimized cell transmission schedule. An
6 example of such a technique can be found in EP-A-702472. This optimized transmission schedule
7 must be recalculated if connections are added or deleted. Consequently, these methods are not easy
8 to implement for use at the very high speeds required by modern networks.

9 **Summary of the Invention**

10 It is an object of the present invention to provide a simple traffic shaping and/or policing
11 technique which can be implemented in hardware at reduced cost.

12 In brief, this object is achieved by a traffic shaping apparatus for packet data communications.
13 The apparatus comprises one or more packet queues for traffic having a plurality of different desired
14 packet transfer rates, one queue being assigned to each desired packet transfer rates. Cell receiving
15 logic is provided to receive incoming packets and to place each incoming data packet in a queue
16 which handles the packet transfer rate for the traffic to which the packet belongs. A plurality of
17 timing circuits are provided, each of which is arranged to operate at a different frequency. The
18 frequencies are selected so that the desired packet transfer rates can be approximated by combining
19 subsets of the available frequencies. The combined outputs of timing circuits in each subset provide
20 packet transfers signals for a particular queue. Cell transmission logic responds to the packet
21 transfer signals to transfer a packet from the corresponding queue to an output.

22 The invention is particularly, but not exclusively, applicable to ATM networks which employ
23 short fixed length cells and in which traffic is carried over a number of connections each having an
24 associated nominal packet transfer rate. In this environment, the system may have one queue for

1 each connection or, alternatively, one queue corresponding to a group of connections having an
2 associated maximum aggregate packet transfer rate. Incoming packets are received and placed in
3 the queue corresponding to the connection or group of connections to which the packet belongs.

4 **Brief Description of the Drawings**

5 The foregoing and other objects, aspects and advantages of the invention may be better
6 understood from the following detailed description of a preferred embodiment of the invention with
7 reference to the drawings, in which:

8 Figure 1 is a schematic diagram showing traffic shaping apparatus;

9 Figure 2 shows the timing unit of the apparatus of Figure 1;

10 Figure 3 is a flow diagram illustrating the operation of the timing unit;

11 Figures 4A, 4B and 4C are timing diagrams illustrating the operation of the timing unit; and

12 Figure 5 shows a switching hub including the traffic shaping apparatus of Figure 1.

13 **Detailed Description of Preferred Embodiment of the Invention**

14 Figure 1 shows in schematic form a traffic shaping device 100 for packet data
15 communications. The apparatus includes a set of cell queues, such as queue 110, which are defined
16 in known fashion using appropriate pointers in a cell buffer memory (not shown). Each queue
17 corresponds either to a particular reserved bandwidth ATM connection having a particular maximum
18 cell transfer rate which was negotiated when the connection was established or to a group of
19 connections for which there is a maximum permitted aggregate bandwidth. Enqueuing logic 120
20 is provided to receive incoming cells from input 125 and place the incoming data cells in the queue

1 assigned to the connection to which the data cells belong.

2 Each queue 110 has a threshold 112 which is set according to the maximum Cell Delay
3 Variation (CDV) allowed for the connection. If the queue is filled to the threshold, any further
4 packets arriving for that queue are discarded by sending them to a free queue 115 which holds all
5 the cells not currently in use by the system. This provides a simple policing function to ensure that
6 the negotiated traffic parameters, in particular the CDV, is not exceeded.

7 A timing unit 130 is provided which includes a plurality of timing circuits 140. For ease of
8 understanding, timing circuits 140 are shown in schematic form only in Figure 1 and will be
9 described in more detail below. As will be understood from the description which follows, timing
10 circuits 140 need not be physically separate within timing unit 130, but may be implemented using
11 a common counter and control logic.

12 Timing circuits 140 are each arranged to operate at a different one of a series of frequencies,
13 $F/2^n$ with $n = 1, 2, 3, \dots, n_{max}$, where F is the maximum packet transfer rate. The effective packet
14 transfer rate for each connection is established by combining outputs from the different timing
15 circuits; that is by creating subsets of packet transfer signals at the different frequencies providing
16 by the timing circuits. Figure 1 shows timing circuits capable of generating packet transfer signals
17 at frequencies $F/2, F/4, F/8, F/16$, and $F/32$. Using timing circuits such as those illustrated, any
18 frequency can be approximated to an accuracy of $F/32$. For example, if a desired cell transfer rate
19 for a particular connection is $15F/32$, this can be represented by $F/4 + F/8 + F/16 + F/32$. In a
20 preferred embodiment, F is 1.5 Mcells/sec which corresponds to a throughput of 636 Mbits/sec for
21 standard fifty-three byte ATM cells..

22 The timing circuits whose outputs are combined to generate a desired packet transfer rate for
23 a particular queue are arranged as described below to generate packet transfer signals on line 145
24 for that queue. As an example, if one of queues 110 is assigned to a connection having a negotiated
25 maximum packet transfer rate of $15F/32$, then outputs of timing circuits capable of generating packet

1 transfer signals at the respective rates of F/4, F/8, F/16 and F/32 would be combined to establish
2 an effective total packet transfer rate of 15F/32 for the connection.

3 Dequeueing logic 150 is responsive to the packet transfer signals to transfer a packet from its
4 queue to output 155.

5 Figure 2 provides more detail about a timing unit 130. The unit 130 comprises a counter 200
6 which is incremented at frequency F. The count value generated by counter 200 is decoded by
7 decoder 210 to produce activating signals on lines 220 at appropriate times. Each of lines 220
8 corresponds to a particular frequency in the series $F/2^n$ with $n = 1, 2, 3, \dots, n_{max}$. Only 5 frequencies
9 are shown in Figure 2, although in practical implementations more frequencies could be used. In
10 a preferred embodiment, a total of 15 different packet transfer frequencies would be generated,
11 which means that any desired packet transfer output can be approximated with a precision of $F/2^{15}$.
12 The activating signals are produced at the start of each period for the particular frequency.

13 There is a phase difference of half a period between the activating signals for neighboring
14 frequencies in the series, so that, for instance, the activating signal for frequency F/4 occurs in the
15 middle of the period for frequency F/2. This spreads the packet transfer signals for each connection
16 over time and thereby reduces the burstiness of the output traffic for each connection.

17 Lines 220 are connected to set/reset circuits 230. There is one set/reset circuit 230 for each
18 frequency in the series. The set/reset circuit for each frequency is set at the beginning of the
19 corresponding period by the signals from decoder 210 on lines 220.

20 Timing unit 130 also includes a memory 250 and a control unit 240, which is linked to each
21 of set/reset circuits 230 via lines 241 and 242.

22 Each frequency of the series has an associated list of queue IDs, represented at 251 in Figure
23 2, stored in memory 250. Also stored in memory 250 is a list 252 of pointers to a current position

1 in each of lists 251. The list of pointers 252 could equally be stored in a suitable register outside
2 memory 250.

3 A simple mechanism, which for clarity is not shown in Figure 2, can be used to prevent
4 setting of any set/reset circuit associated with a frequency not currently needed. In a preferred
5 embodiment, this mechanism includes a register which is arranged to keep a record of which
6 frequencies are not currently needed; i.e., which have no queues in their associated list 251. Lines
7 220 are gated to set/reset circuits 230 under the control of this register. If packet transfer signals at
8 a particular frequency are not needed, as indicated by the absence of a queue identifier in the list 251
9 associated with the particular frequency, then decoder 210 is not allowed to set the corresponding
10 set/reset circuit 230.

11 Control unit 240 is arranged to poll, over lines 242, the state of set/reset circuits 230 at
12 frequency F and generate on line 145 a signal identifying a particular cell queue using the lists stored
13 in memory 250. A new cell queue is thus identified at frequency F and consequently the rate at
14 which cells are passed to output 155 from queues 110 is equal to frequency F.

15 This process is illustrated in Figure 3 in the form of a simple flow diagram. Set/reset circuits
16 230 are polled in order of decreasing frequency to find the highest frequency for which the respective
17 set/reset circuit is set. This is represented by step 300. For the first set/reset circuit which is set,
18 control unit 240 identifies the current queue ID from a list of queues associated with that frequency
19 and outputs the queue ID on line 145. This is represented by step 310.

20 If the queue ID is not last in the list, the pointer is incremented to point to the next queue ID
21 in the list and control returns to step 300. If the queue ID is the last in the list, control unit 240
22 deactivates the set/reset circuit corresponding to that frequency by asserting a reset signal on the
23 appropriate one of lines 241. This is represented by step 320. The pointer to the list 251 for that
24 frequency in the list of pointers 252 is set to point to the first connection ID in the list and control
25 returns to step 300.

1 Set/Reset circuits 230 are "set" dominant so that if a set signal from decoder 210 occurs at
2 the same time as a reset signal from control unit 240, the set/reset circuit remains in the set state.

3 The fact that the set/reset circuits are polled in order of decreasing frequency provides a
4 priority control mechanism which ensures that the timing circuits can only generate a packet transfer
5 signal if there is no timing circuit of higher frequency wishing to generate a packet transfer signal
6 at that time. In this way, connections having high data transfer rates are given higher priority since
7 the impact of a given cell delay is a direct function of data transfer rates.

8 Figures 4A, 4B and 4C represent signals generated in applying this process. Dotted vertical
9 lines, such as line 405, indicate the beginning of each time period for each of the different
10 frequencies, illustrating the phase difference introduced between neighboring frequencies in the
11 series. These dotted lines correspond to the signals produced by decoder 210 on the lines 220.

12 In Figure 4A, connection identifiers asserted on line 145 at given times are shown in the
13 OUT line at the top of the Figure. The other solid lines of Figure 4A show the state of the set/reset
14 circuits 230 corresponding to the frequencies listed on the left hand side of Figure 4A.

15 In Figure 4A, the scheduling of three different connections is illustrated. The connections
16 are identified by the connection identifiers 1, 2 and 3 respectively, in the OUT line. Connection 1
17 requires a packet transfer rate of F/2. Connection 2 requires a packet transfer rate of 9F/32 or F/4
18 + F/32. Connection 3 requires a packet transfer rate of 3F/32 or F/16 + F/32. Consequently, the
19 contents of lists 251 would, in this example, be as follows: the list for frequency F/2 would contain
20 connection ID 1; the list for frequency F/4 would contain connection ID 2; the list for frequency F/8
21 would be empty; the list for frequency F/16 would contain connection ID 3; and the list for
22 frequency F/32 would contain connection IDs 2 and 3.

23 Once the list of connection IDs for a particular one of the frequencies has been exhausted,
24 the corresponding set/reset circuit 220 is reset via a signal on the appropriate one of lines 241. This

1 is illustrated by the downward edges, such as downward edge 410 in Figures 4A-C. It can be seen
2 in Figure 4A that the set/reset circuit for frequency F/32 remains active for several cycles until the
3 ID for both connections 2 and 3 have been asserted. Only then is the set/reset circuit for frequency
4 F/32 deactivated.

5 In Figure 4A it can be observed that, for instance, connection 2 is asserted both when
6 frequency F/4 is the highest active frequency, as illustrated at 415, and also when F/32 is the highest
7 frequency active, as illustrated at 420. This is due to the presence of connection id 2 in the lists 251
8 corresponding to both these frequencies.

9 In Figure 4B, connection 1 has been eliminated and a new connection 4 added with a
10 required packet transfer rate of 12F/32 or F/4 + F/8. Consequently, the contents of lists 251 are
11 modified as follows: the list for frequency F/2 is empty, connection 1 having been removed; the list
12 for frequency F/4 contains connection ID 2 and connection ID 4; the list for frequency F/8 contains
13 connection ID 4; the list for frequency F/16 still contains connection ID 3; and the list for frequency
14 F/32 still contains connection IDs 2 and 3.

15 As before, the connection identifier asserted on line 145 at any particular time is shown in
16 the OUT line of Figure 4B. The other lines of Figure 4B show the state of the set/reset circuits 230
17 corresponding to the frequencies listed on the left hand side of Figure 4B.

18 In Figure 4C, connection 2 has been eliminated and new connections 5, 6 and 7 added.
19 Connections 5 and 7 each have a packet transfer rate of 3F/32 or F/16 + F/32. Connection 6 has a
20 desired packet transfer rate of F/32. Consequently, the contents of lists 251 are modified as follows:
21 the list for frequency F/2 is still empty; the list for frequency F/4 contains only connection ID 4; the
22 list for frequency F/8 still contains connection ID 4; the list for frequency F/16 contains connection
23 IDs 3, 5 and 7; and the list for frequency F/32 contains connection IDs 3, 5, 6 and 7.

24 As before, the connection identifier asserted on line 145 at any particular time is shown in

1 the OUT line of Figure 4C. The other lines of Figure 4C show the state of the set/reset circuits 230
2 corresponding to the frequencies listed on the left hand side of Figure 4C.

3 New connections or groups of connections corresponding to queues 110 can be easily added
4 or deleted merely by creating or deleting a corresponding queue 110 and changing the lists 251 so
5 that a packet transfer signal is produced on line 145 at a corresponding rate.

6 Figure 5 is a schematic diagram showing an ATM switching hub in which such a traffic
7 shaping device is used. The hub comprises a switching fabric 500 and a number of adapters 510 of
8 which only one is shown in Figure 5. Each adapter 510 is divided into two parts, a receive part 510A
9 and a transmit part 510B. ATM cells are received on inputs 520, multiplexed within the receive part
10 510A of an adapter and transmitted via switching fabric 500 to the transmit part 510B of one of the
11 adapters, where they are demultiplexed and output on one of output lines 530.

12 The shaping device 100 is located in the receive part 510A of each adapter and is arranged
13 to shape and police the traffic being passed through switching fabric 500. In the parlance of ATM
14 technology, shaping device 100 in the receive part 510A has the function of UPC (Usage Parameter
15 Control) or NPC (Network Parameter Control).

16 A similar shaping device, denoted 100' in Figure 5, can also be located in the transmit part
17 of the adapter for the purposes of output shaping. In this case, each queue would normally represent
18 a number of connections which are destined for transmission over a single ATM connection at, for
19 example, a LAN/WAN interface. For output shaping, the policing function of the shaping device,
20 ie the thresholds 112, is generally not required because such function is normally implemented at
21 the input of a node.

22 While the invention has been described in terms of preferred embodiments, those skilled in
23 the art will recognize that the invention can be practiced with modifications within the spirit and
24 scope of the appended claims.